How Discourse Adjectives Synchronize the Speaker and the Listener's Beliefs
(*It is clear that Briscoe is a Detective*)

Gina Taranto
H5 Technologies, San Francisco, CA

1. Introduction

This paper presents Discourse Adjectives, a natural class of predicate whose members include *apparent*, *evident*, *clear*, and *obvious*, in their use as propositional modifiers in sentences such as those in (1).

(1)   a. It is *apparent* that somebody committed a crime.
          b. It is *evident* that the police are on their way.
          c. It is *clear* that Briscoe is a detective.
          d. It is *obvious* that someone watches **way** too much Law & Order.

The main claim is that Discourse Adjectives mean in an interesting and important way: their use is entirely metaliguistic. They offer interlocutors a way to talk about their conversation, rather than talking about their world. As such, they appear to violate a commonly held constraint on assertions: they do not appear to add new information to the common ground of a discourse. This is interesting because Discourse Adjectives represent the first identified class of lexical predicates whose use is entirely metalinguistic. Without minor modifications, previous semantic models cannot account for their semantics. In particular, any analysis will need to combine a finely articulated representation of the common ground with a mechanism for acknowledging degrees of probability.

2. Discourse Adjectives are "Factive Impostors"

An initial exploration of the semantics of Discourse Adjectives might hypothesize that they're factive, in the sense of Kiparsky and Kiparsky (1970). Consider (2).

(2)   It is clear that Briscoe is a detective.

Example (2) seems to entail that Briscoe is a detective. However, standard diagnostics for factivity show that predications of clarity do not entail the truth of their complement propositions. The truth of the complement is not an entailment under either negation or questioning, as shown in (3).

(3)   a. It isn't clear that Briscoe is a detective, which makes sense, since he's
           not a detective.
       b. Is it clear that Briscoe is a detective?

(3a) is not contradictory, and (3b) is indeterminate between a reading in
which the speaker questions Briscoe's detectivehood, and one in which she
wonders whether this proposition is 'clear' to the addressee. In spite of the fact
that Discourse Adjectives fail the two standard diagnostics for factivity, it is hard
to deny the strong flavor of factivity that sentence (2) carries. That is, it really
feels like *it is clear that Briscoe is a detective* entails that Briscoe is a detective.

3. A Case Study of the Semantics of *Clear*

3.1. The Paradox of Asserting Clarity

The starting point for an analysis of *clear* is a quick review the Paradox of
Asserting Clarity, first introduced in Barker and Taranto (2002). The dilemma
that underpins the paradox arises from the standard assumption in (4), that an
assertion is felicitous only if it adds new information to the common ground.
After all, what use could it be to claim that a proposition is true if it is already
accepted as true? The answer advocated by Barker and Taranto is given in (5).

(4)   Informativeness Constraint on Assertions
       An assertion is felicitous only if it adds new information to the common
       ground. (after Stalnaker 1979:325, van der Sandt 1992:367)

(5)   Some sentences can have side-effects besides adding new information
       about the world to the common ground, and it can be worth asserting a
       sentence entirely for the sake of these side-effects. (Following Beaver
       (2002))

Returning to the central example for this paper, the sentence in (2), the
observation is that intuitively, if (2) is true, then before it is uttered, the speaker
and the addressee must already believe that Briscoe is a detective: if either is not
already convinced that Briscoe is a detective, then it isn't clear at all. But, if it
was already evident that Briscoe is a detective, then asserting (2) adds no new
information to the context, contra the assumption spelled out in the
Informativeness Constraint in (4). This leads to Lemma 1, elaborated in (6).

(6)   Lemma 1
       If (2) is true, it adds no new information to the context, so why bother to
       assert it?

Lemma 1 could be resolved by showing that an utterance of (2) does add new information to the common ground namely, the information that Briscoe is a detective. However, this would give rise to Lemma 2, which is elaborated in (7).

(7)   Lemma 2
       Assume (2) entails Briscoe is a detective. If the speaker has decided that Briscoe is a detective on the basis of information shared by the speaker and addressee, then (2) only adds information if the speaker assumes that the addressee has not come to the same conclusion, in which case it is *not* in fact clear that Briscoe is a detective, otherwise the addressee would have realized it on her own.

In other words, perhaps (2) can only be used in situations in which the addressee is an idiot: she has all the evidence she needs to realize that Briscoe is a detective, but she fails to take that last logical step. To consider such a situation, assume that the propositions in (8) are true, Briscoe, who works for the NYPD detective division, is walking alongside Detective Logan, and he's waving his pistol, badge, and magnifying glass.

(8)   a.   Briscoe works for the NYPD detective division.
       b.   Briscoe's partner is a detective, Detective Logan.
       c.   Briscoe carries a pistol and magnifying glass wherever he goes.
       d.   Briscoe carries a detective's badge.

To be even more explicit, imagine that the sentences in (8) have just been uttered by B. That's when A replies with her own version of (2), an utterance of the sentence in (9).

(9)   Clearly you fool, Briscoe is a detective!

3.2. The Framework - Gunlogson (2001)

The starting point for the formalization of the proposed analysis of the semantics of *clear* is Gunlogson's adaptation of a Standard Stalnakerian Model. Some of the key points of this model are outlined in (10).

(10)   Gunlogson's Adaptation of The Standard Model
       a.   Within the Common Ground a distinction is made between the
            Discourse Commitments of the speaker and the addressee.
       b.   An individual's Discourse Commitments are identified with that
            individual's public beliefs.

   c. The Context Set of a Discourse with two participants A and B
       $(CS_{A,B})$ is an ordered pair of sets of the discourse commitments of A
       and B: the commitment set of A (cs_A) includes the worlds of which
       A's public beliefs are true; the commitment set of B (cs_B) includes the
       worlds of which B's public beliefs are true.
   d. Asserting a sentence S: remove all worlds from the commitment set of
       the speaker in which S is not true.

   Points in (10a through c) establish that the common ground of a discourse
is construed as the individual sets of possible worlds corresponding to the way the
discourse participants individually believe the world might be.   Point (10d)
explains that when a sentence is asserted the common ground is updated by
removing from the speaker's commitment set all of the worlds in which the
proposition expressed by the sentence is not true. An example of the affect of an
utterance of *Briscoe is a detective* is provided in (12).   The model of possible
worlds under consideration is in (11). Odd numbered worlds are worlds in which
Briscoe is a detective, and this is emphasized with bold-type.   Even numbered
worlds are worlds in which Briscoe is not a detective.

(11)   Model of Possible Worlds

       w_9 =   **Briscoe is a detective**
       w_{10} =   Briscoe is not a detective
       w_{11} =   **Briscoe is a detective**
       w_{12} =   Briscoe is not a detective

(12) cs_A cs_B
    a. Input Context C
                              { w_9, w_{10}, w_{11}, w_{12} }   { w_9, w_{10}, w_{11}, w_{12} }
    b. C + ⟦ Briscoe is a detective⟧
                              { w_9, w_{11}, }   { w_9, w_{10}, w_{11}, w_{12} }

   The context depicted in (12a) is neutral toward the proposition expressed
by *Briscoe is a detective*. The individual commitment sets of both participants
include worlds in which Briscoe is, as well as worlds in which Briscoe is not, a
detective. Update with A's assertion of *Briscoe is a detective* removes worlds 10
and 12, worlds in which Briscoe is not a detective, from A's individual
commitment set. The resulting context in (12b) is a context that is biased toward
the truth of the proposition expressed by *Briscoe is a detective*, since it is a public
commitment of A, but not B, in line with the definition in (14). Gunlogson's
definitions for *bias* and *neutrality* are provided in (13) and (14).

(13)   C is *biased toward* p iff p is a public commitment of exactly one discourse participant.

(14)   C is *neutral toward* p iff neither discourse participant is committed to either p or W-p.

3.3. Resolving the Paradox

With the basics of Gunlogson's adaptation of a standard Stalnakerian model established, it is possible to continue an analysis of the semantics of *clear* by picking up with the with the observation that Discourse Adjectives aren't factive.   The proposed analysis respects the fact that there are a number of possibilities for the way the world might be, and that these are ordered in terms of their plausibility.   For instance, regarding the proposition that Briscoe is a detective, in any given world the possible explanations might include that Briscoe is a detective, as in world nine, or that he's not a *real* detective, he's getting dressed up for a Halloween party, or perhaps even the more unlikely alternative in world twelve, that Briscoe is not a detective at all, but the CIA wants us to believe that he is.

(15)   Possible explanations for the evidence suggesting that Briscoe is a detective, from most to least likely:

w_9 = **Briscoe is a detective.**
w_{10} =   Briscoe isn't a real detective, he's getting ready for a Halloween party.
w_{11} = **The pistol and magnifying glass belong to Briscoe's brother, but coincidentally Briscoe is a detective, too.**
w_{12} = Briscoe isn't a detective, but the CIA wants us to believe that he is.

Nothing in the sentences in (8) rules out any of the possible worlds given in (15).   But because in some of these worlds Briscoe isn't a detective, update with sentence (2), *it is clear that Briscoe is a detective* will at least eliminate those worlds in which Briscoe isn't a detective. This result is suspicious, since simply asserting that *Briscoe is a detective* will achieve the same result. This leads to the following question: why not just assert that Briscoe is a detective, why ever assert that it's *clear* that Briscoe is a detective? At this point it should be clear why a speaker might be reluctant to assert that Briscoe is a detective: because Briscoe might not be a detective – there are other live possibilities.

Since it is known from Grice that it is uncooperative to assert something for which one lacks sufficient evidence, it would be uncooperative to claim that Briscoe is a detective unless one was absolutely certain. If this thinking is on the

right track, the conclusion is that clarity is asserted only in contexts in which there is some lingering uncertainty that the complement is in fact true. However, if this is right, it leads to a new problem: the original paradox has been reconstructed, only in reverse. The reconstructed paradox is given in (16).

(16)   The Reconstructed Paradox
        It is clear that p is only asserted in situations in which it is in fact **not** clear that p!

        The key to resolving this paradox depends on appreciating how the grammar deals with degrees of probability. In other words, the appropriateness of asserting clarity depends on the degrees of probability of different explanations for the facts. Situations in which the applicability of a predicate depends on degrees are well known in the literature: what we're dealing with here is vagueness, or the observation that in a given situation, it's not always clear *how clear* a proposition needs to be to count as *clear*.

3.3.1. Vagueness (Fine 1975, Williamson 1994, 1999, Kennedy 1997)

Following Barker (2002), and Barker and Taranto (2002) the proposed analysis makes use of a delineation function that takes a situation and an adjective meaning, and returns the vague standard for the adjective in the given situation. This is written as in (17).

(17)   **d**(c)([[*clear*]])
        d: delineation function (Lewis 1970): takes·a situation *c* and an adjective meaning and returns the vague standard for that adjective in the given situation.

        The idea is that different possible worlds may differ precisely in the standard they impose on a vague adjective. So in one world the absolute standard may be set higher or lower than it is in another world. By using a mechanism that allows for a representation of the vagueness inherent in predications of clarity, the truth conditions of (18a) *That Briscoe is a detective is clear* can be characterized as in (18b).

(18)   a.   That Briscoe is a detective is clear.
        b.   The maximal degree to which the proposition expressed by *Briscoe is a detective* is clear is *at least as great* as **d**(c)([[*clear*]])

        (18b) merely states that the maximal degree to which the proposition expressed by *Briscoe is a detective* is clear is at least as great as the value returned by the delineation function. The only information provided by an utterance of

(18a) regards the prevailing standard for clarity. That is, information is provided about the state of the discourse, but no new information is provided about the facts of the world.

It is easy to model the use of vague predicates in a Stalnakerian model of context-update if as long as a few natural assumptions are made. The necessary assumptions are outlined in (19).

(19)   During a conversation, one thing is certain about the world; that a conversation is taking place, that the speaker is speaking, that the hearer is being addressed, and so on. Therefore every possible world in the initial context will be a world in which the conversation underway is taking place.

When the assumptions in (19) are incorporated with what is known about vague predicates the result is the observation in (20): that worlds in a context may vary with respect to the values assigned by delineation functions.

(20)   One way in which the worlds in a context may vary is in the value of their delineation function for the version of the conversation in that world.

Based on the assumptions of (19) and (20), a first attempt at an analysis of *clear* is presented in (21).

(21) $cs_x + [\![$ *it is clear that p* $]\!] = \{$ $w \in cs_x$: the maximal degree to which p is likely to be true in cs_x is at least as great as $\mathbf{d}(c)([\![$ *clear* $]\!]$) in $w\}$ [to be revised in (24)]

The proposal in (21) claims that updating the commitment set of a discourse participant x with the information expressed by *it is clear that p* will yield a set of worlds in which the maximal degree to which p is likely to be true in x's commitment set is as least as great as the delineation function requires in that world. This analysis captures the connection between likelihood and clarity, and specifies the respect in which asserting clarity is similar to asserting the applicability of a vague predicate. However, it cannot be right.

The problem is that in Gunlogson's model (as well as other Stalnakerian models), propositions don't have probabilities. For any given possible world, either Briscoe is a detective, or he's not. No matter what the standard for clarity is, worlds in which the probability is 1 will survive update according to (21), and worlds in which the probability is 0 will not. Thus, in the absence of acknowledging probability, the meaning of *it is clear that Briscoe is a detective* is identical to the meaning of *Briscoe is a detective,* which has already been shown to be incorrect.

3.4. The Solution

This problem raised by the analysis in (21) can be solved by building on the observation that likelihood is a judgment made by some sentient creature who is contemplating p. Therefore, if likelihood plays a role in assertions of clarity, we need to figure out who is judging likelihood. An important clue to the identity of the experiencers of clarity comes from comparing a simple assertion of clarity to one in which the experiencer is overt. Consider the data in (22).

(22) a. It is clear that Briscoe is a detective.
          ASSERTION OF SIMPLE CLARITY
      b. It is clear to me that Briscoe is a detective.
          ASSERTION OF PERSONAL CLARITY
      c. (Surely) It is clear to you that Briscoe is a detective.
          ASSERTION OF PERSONAL CLARITY

In the revised analysis, the meaning of (22a) is approximated as the conjunction of (22b and c): if it is clear that Briscoe is a detective, then it is clear to the discourse participants that Briscoe is a detective. This suggests refining the context change potential in (21) as in (23), which considers judgments of likelihood at each world. So (23) recognizes that belief is a gradient attitude, and behaves just like any other vague predicate.

(23) $cs_x + [\![$ *it is clear that p* $]\!] = \{ w \in cs_x :$ the maximal degree to which the counterpart of x judges that p is likely to be true in cs_x is at least as great as $\mathbf{d}(c)([\![clear]\!])$ in $w\}$

In practical terms, this means that if a speaker asserts (24a) *it is clear to me that Briscoe is a detective* with the semantics in (24b), then only those worlds will survive update at which the speaker *believes* that Briscoe is a detective.

(24) a. It is clear to me that Briscoe is a detective.
      b. $\{w$ such that $\mathbf{believe}(w)(\mathbf{d},\text{speaker},\text{Briscoe-is-a-detective}) \geq \mathbf{d}(c)([\![clear]\!])\}$

Worlds that are excluded will include worlds in which there is enough uncertainty to reduce the speaker's belief of Briscoe's detectivehood to below that world's specified threshold for clarity. Worlds may survive in which Briscoe is not a detective, as long as the speaker believes that Briscoe is a detective in that world.

An example is provided in (26), which includes information about Briscoe's being a detective, as well as information about the clarity of this

proposition in terms of the discourse participant's beliefs. In (25), "D" abbreviates the proposition expressed by *Briscoe is a detective*, and "Bel" is a belief operator used to indicate whether A or B believe that Briscoe is a detective. In (25), world nine is a world in which both A and B believe Briscoe is a detective, and world ten is a world in which neither A nor B believe that Briscoe is a detective.

(25)   Expanded Model of possible worlds

w_9	=	**Briscoe is a detective**	Bel(A,D)	Bel(B,D)
w_{10}	=	Briscoe is not a detective	¬Bel (A,D)	¬Bel (B,D)
		(Halloween Party)		
w_{11}	=	**Briscoe is a detective**	Bel(A,D)	¬Bel (B,D)
w_{12}	=	Briscoe is not a detective	Bel(A,D)	¬Bel (B,D)
		(CIA plot)		

Imagine a situation in which the speaker, A, is unaware of the CIA plot; her utterance of (24a) is modeled in (26). The input context is as modeled in (26a), and the result of her utterance eliminates world ten – the world in which A does not believe Briscoe is a detective – from her commitment set. Her commitment set still includes world twelve, a world in which Briscoe isn't a detective, but because in this world there is sufficient evidence to persuade our speaker that Briscoe is a detective, this world remains a live possibility.

(26)
    a. Input Context C $\{ w_9, w_{10}, w_{11}, w_{12} \}$ $\{ w_9, w_{10}, w_{11}, w_{12} \}$
    b. C + A: ⟦(26a)⟧ $\{ w_9, w_{11}, w_{12} \}$ $\{ w_9, w_{10}, w_{11}, w_{12} \}$

with column headers cs_A and cs_B.

In a departure from Gunlogson (which will be addressed below), the analysis of *clear* presented here adopts Walker's Collaborative principle, shown in (27), which holds that conversants must provide evidence of discrepancy in belief as soon as possible.

(27)   Collaborative Principle (Walker 1992)
        Conversants must provide evidence of a discrepancy in belief as soon as possible.

The claim made here about the semantics of *clear* is that in the normal course of events, when A utters *it is clear to me that Briscoe is a detective*, B will have no choice but to accept the fact that it *is* clear to A – A is the highest authority on A's beliefs. By brining in Walker's Collaborative Principle, it is possible to formalize what happens if B does not immediately express doubt about the truth or sincerity of A's statement. If B remains silent, then the discourse

model will reflect individual public commitments on the part of both A and B to A's belief in Briscoe's detectivehood. Thus, the representation of B's commitment set as depicted in (26b) is incomplete – B's commitment set must also reflect her belief (- or acquiescence -) to A's commitment to the clarity of the proposition expressed by *Briscoe is a detective*, as shown in (28).

(28) cs_A cs_B

     a.   Input Context C $\{ w_9, w_{10}, w_{11}, w_{12} \}$ $\{ w_9, w_{10}, w_{11}, w_{12} \}$
     b.   C + A: $[\![(26a)]\!]$ $\{ w_9, w_{11}, w_{12} \}$ $\{ w_9, w_{10}, w_{11}, w_{12} \}$
     c.   No objection from $\{ w_9, w_{11}, w_{12} \}$ $\{ w_9, w_{11}, w_{12} \}$

Note that the commitment set of B in (28c) includes two worlds in which it is not clear to B that Briscoe is a detective – Her commitment set includes worlds eleven and world twelve, worlds in which it is clear to A that Briscoe is a detective, even though it is not clear to B that Briscoe is a detective.

     Example (29) shows the update effect of an assertion of ***simple clarity***.

(29)   It is clear that Briscoe is a detective.

(30) cs_A cs_B

     a.   Input Context C $\{ w_9, w_{10}, w_{11}, w_{12} \}$ $\{ w_9, w_{10}, w_{11}, w_{12} \}$
     b.   C + A: "(31)" $\{ w_9 \}$ $\{ w_9, w_{10}, w_{11}, w_{12} \}$
     c.   No objection from B $\{ w_9 \}$ $\{ w_9 \}$

Since the semantics of *clear* specify that the default interpretation of the experiencer is as the Discourse Participants, the only world that survives update is world nine, the sole world in which it is clear to both A and B that Briscoe is a detective. By this analysis dialogs involving the denial of simple assertions involve contradiction and repair, while denials of assertions of personal clarity do not. Consider (31).

(31)   <u>Denial of simple assertions involve contradiction and repair</u>
       A: Briscoe is a detective.
       B: Actually, Briscoe isn't a detective. I just asked him and he proved
           he's not.

In (31), a dialog in which A utters "Briscoe is a detective", and B replies "Actually, he's not a detective. I just asked him and he proved he's not" B's statement *contradicts* A's. Presumably, some form of repair must occur before the conversation can proceed. In contrast, the dialog in (32) includes A's assertion of personal clarity.

(32)   Denial of assertions of personal clarity do not involve contradiction
        A: It is clear to me that Briscoe is a detective.
        B: Actually, Briscoe isn't a detective. I happen to be a CIA operative,
            and I can tell you that it's all part of a supremely devious charade.

When A utters "It is clear to me that Briscoe is a detective", and B responds "Actually, Briscoe isn't a detective. I happen to be a CIA operative, and I can tell you that it is all part of a supremely devious charade", B has not contradicted A: it remains true that it was clear to A that Briscoe was a detective. Whether or not A will continue to believe in Briscoe's detectivehood will depend on her belief in the truth of B's statement.

3.4.1. Explaining "Factive Impostors"

It is now possible to explain the status of Discourse Adjectives as "Factive Impostors". That is, it is possible to explain the observation made in the introduction that in spite of the fact that Discourse Adjectives fail the standard diagnostics for factivity, they intuitively "feel" very factive. The explanation available given the analysis presented here is that in the absence of an overt experiencer, the entities doing the experiencing default to the speaker and addressee. The semantics of simple clarity guarantee that in the absence of an immediate and overt objection by the addressee, *every world* in the updated context will be a world in which the experiencer – in this case, both the speaker and the addressee - *believes* the truth of the proposition. The result is that, in terms of the model relevant to a particular discourse, the discourse participants believe the proposition is likely to be true in every world in the updated context. As spelled out in (33), asserting *It is clear that p* does not entail *p*, but it guarantees that the discourse participants are *justified* in behaving as if *p* is true.

(33)   •   In the absence of an overt experiencer, the entities doing the
            believing default to both the speaker and the addressee.
        •   The semantics of clarity guarantee that every world in the updated
            context will be a world in which the experiencer believes the truth of
            the proposition.
        •   The result is that all of the discourse participants believe the truth of
            the proposition in every world in the updated context.

The key thing to note is that asserting clarity is purely about the judgment of the discourse participants, and not about what is the case in any part of the world under discussion. This is precisely how asserting clarity synchronizes the common ground: it forces the speaker and addressee to acknowledge that they are in a position to treat a certain proposition as if it were a fact.

(34)   Utterances with Discourse Adjectives can help to synchronize the Context
        Set of a discourse by elevating the shared private beliefs of the discourse
        participants to the status of mutual public beliefs.

        This explanation of Factive Impostors provides a resolution to Lemma 2 -
an assertion of personal clarity doesn't require that the speaker believes the
addressee is an idiot, it merely requires that she believes a private belief is shared,
and could be made public.

(35)   Resolution of Lemma 2
        Isn't asserting clarity tantamount to suggesting that the hearer is an idiot?
        Resolution: It is sufficient for the speaker to assume that the addressee
        may not know that the speaker also believes that Briscoe is a detective.

        As for Lemma 1, if asserting clarity adds no new information about the
situation under discussion, what use is it to assert it.   The resolution is that
information is added.   The new information regards the attitude of the discourse
participants *towards* the proposition in question.

(36)   Resolution of Lemma 1
        If asserting clarity adds no new information about the situation under
        discussion, what use is it to assert it?
        Resolution: information *is* added about the attitude of the discourse
        participants.

        As for the reconstructed paradox, that *it is clear that p* is asserted only in
situations in which it is in fact not clear that *p*, the resolution comes from the
understanding that asserting clarity does not require asserting perfect clarity. By
recognizing the role of vagueness, it is possible to capture how asserting clarity
means that the proposition in question is merely clear enough – in particular, we
see that asserting clarity means that the proposition is clear enough to proceed as
if it were true.

(37)   Resolution of The Reconstructed Paradox
        It is clear that p is only asserted in situations in which it is in fact **not** clear
        that p.

4. Notes on the Implementation of Gunlogson: Degrees of Commitment

The final section of this paper will briefly discuss a consequence of updating
Gunlogson's model with degrees.   This modification provides a means to
represent degrees of public commitment to a proposition, which is good, since a

comparison across the class of Discourse Adjective reveals that an individual's degree of public commitment *toward* a *proposition* can be vague.

Consider the sentences in (38), as possibly uttered by Briscoe upon his arrival and brief survey of the scene of the crime.

(38)   a. It is apparent that Mr. Body was murdered in the conservatory with the candlestick.

          b. It is evident that Mr. Body was murdered in the conservatory with the candlestick.

          c. It is clear that Mr. Body was murdered in the conservatory with the candlestick.

          d. It is obvious that Mr. Body was murdered in the conservatory with the candlestick.

Imagine that the truth is that Mr. Body was murdered in the Billiard Room with the lead pipe. Intuition suggests us that if Briscoe utters (38a) and later finds out that he was wrong, we might *reasonably* expect him to be mildly surprised. However, if Briscoe had instead uttered (38d), an observer might reasonably find it odd if he showed only mild surprise when he learned that his analysis was incorrect. This suggests that "obvious" expresses a higher degree of public commitment to a proposition than does "apparent".

Consider the data that drove Gunlogson to isolate individual commitment sets. Her concern was the data shown in (39): the declarative sentence with falling intonation *it's raining*, the declarative sentence with rising intonation *it's raining?*, and the interrogative sentence with rising intonation *is it raining?*

(39)   a. It's raining.   DECLARATIVE FORM, FALLING INTONATION.

          b. It's raining?   DECLARATIVE FORM, RISING INTONATION.

          c. Is it raining?   INTERROGATIVE FORM, RISING INTONATION.

Gunlogson notes that with falling intonation, an utterance of *it's raining* implicates a *speaker's* commitment to the content of the proposition expressed by that sentence, while an utterance of *it's raining?*, with rising intonation, implicates commitment on the part of the addressee. This is why, for example, an utterance like (40) has what might be considered an impolite implication: *Ah, so you've stopped drowning kittens?* suggests that the addressee has in fact previously drowned kittens.

(40)   Ah, so you've stopped drowning kittens?

A significant modification that I have made to Gunlogson's model is that I allow utterances to *directly* influence the commitment sets of individual discourse participants. Gunlogson merely wanted access to individual commitments to

define the notion of Contextual Bias, which allowed her to state a condition on the felicitous utterance of a rising declarative. She notes that a condition on the appropriate use of (39b) or (40) is that the context must already be biased toward the propositions expressed by *it's raining* or *you've stopped drowning kittens* by virtue of the addressee's prior public commitment to them. She achieves this by positing the Contextual Bias Condition, stated descriptively in (41).

(41)   Contextual Bias Condition (Gunlogson 2001:6)
        Rising declaratives can only be used as questions in contexts where the
        Addressee is already publicly committed to the proposition expressed.

The analysis presented here is not meant to dispute Gunlogson's analysis. The claim is merely that an accurate description of the semantics of Discourse Adjectives require more than the notion of contextual bias. Gunlogson shows that rising declaratives pattern like interrogatives in allowing a reading in which the speaker is understood to be skeptical of the proposition expressed. Falling declaratives cannot co-occur with overt expressions of skepticism. An example she provides to illustrate this fact is given in (42).

(42)   [A & B are looking at a co-worker's much-dented car]
        A: His driving has gotten a lot better.
        B's response:
            a. Has it? I don't see much evidence of that.
            b. It has? I don't see much evidence of that.
            c. It has. #I don't see much evidence of that.

In terms that incorporate vagueness, this can be restated as follows: rising intonation signals a degree of commitment to the truth of a proposition that is *less than* the minimum standard for absolute commitment, while falling intonation signals a degree of commitment that is *at least as great as* that standard. This can be translated into a diagnostic for degrees of commitment involved with derivationally related Discourse Adverbs, as shown in (43).

(43)   [Office Gossip maintains that a co-worker's driving has recently
        improved. A & B are looking at that co-worker's much-dented car, A says
        to B]:
            a. Apparently her driving has gotten a lot better, but I'm not so sure
                it's true.
            b. Evidently her driving has gotten a lot better, but I'm not so sure
                it's true.
            c. #Clearly her driving has gotten a lot better, but I'm not so sure it's
                true.

d. #Obviously her driving has gotten a lot better, but I'm not so sure it's true.

The difference in acceptability of these utterances is explained through consideration of degrees of probability. Specifically, the adverbs derived from *apparent* and *evident* impose relatively loose standards with respect to the minimum degree of probability they tolerate. This allows the speaker to remain skeptical of the truth of the proposition expressed by *her driving has gotten a lot better*. Because of this, the assertions with *apparently* or *evidently* are compatible with overt expressions of doubt, as in (43a,b).

In contrast, the sentences with *clearly* and *obviously*, which impose a higher minimum standard of probability, and do not allow such an overt expression of doubt. Inconsistency results when sentences with adverbs derived from *clear* and *obvious* are uttered with a skeptical follow-up. Thus, *clear* and *obvious* pattern like falling declaratives in strongly committing the speaker to the proposition expressed.

These points about degrees of commitment lead to the final point to make regarding the notion of commitment. This point has to do with the attribution of commitment to individual participants in a discourse. As summarized in (44), Gunlogson shows that a speaker can use falling intonation to signal her own commitment to the propositional content of an utterance, and a speaker can use RISING intonation to signal commitment on the part of the **addressee**.

(44)   Attribution of Commitment in Gunlogson's Model
    a. The form of a sentence compositionally contributes commitment.
        • DECLARATIVE SENTENCE FORM contributes commitment
        • INTERROGATIVE SENTENCE FORM contributes a lack of commitment
    b. The intonational contour of a sentence compositionally attributes commitment to a discourse participant.
        • FALLING INTONATION attributes commitment to the speaker
        • RISING INTONATION attributes commitment to the addressee

A natural question to ask following this is: how does a speaker signal commitment on the part of both herself and the addressee – that is, all of the discourse participants – to the propositional content of a declarative sentence? There appears to be no intonational contour that serves this function in English. But the analysis presented in this paper shows that Discourse Adjectives fill this role in the grammar. The use of a Discourse Adjective in a declarative sentence with falling intonation is a strategy a speaker can adopt to commit both speaker and addressee to the content of a proposition by uttering a single sentence, such as the central example of this paper, repeated in (45).

(45)   It's clear that Briscoe is a detective.

Further, this analysis of Discourse Adjectives combined with Gunlogson's compositional analysis of rising intonation makes a prediction about the meaning of an utterance of a declarative sentence with a Discourse Adjective and rising intonation, as in (46).

(46)   It's clear that Briscoe is a detective?

The prediction is that an utterance of (46) signals a commitment on the part of the addressee to the belief that the interlocutors agree that Briscoe is a detective. This prediction is borne out. The combination of Gunlogson's Cntextual Bias Condition with the semantics provided here for Discourse Adjectives, accurately captures the semantics of (46), as well as the intuition that its utterance can only be felicitously uttered in a situation that is contextually biased towards both discourse participants already being publicly committed to the proposition expressed by (45).

5. In Conclusion

This paper has introduced Discourse adjectives, explained their ability to seem factive, and in so doing has shown that when an assertion of simple clarity is accepted in the discourse model, the interlocutors are licensed to proceed as if the designated proposition is true, if only for the sake of the current discussion. The analysis provides empirical support for Stalnaker's claim that the discourse model must contain a representation of the conversation taking place, and Gunlogson's call for a finely articulated model of the Common Ground that distinguishes between the public commitments of the speaker and addressee. It should be noted, however, that the implementation presented here may go beyond what either of these researchers originally intended. Finally, as for Discourse Adjectives themselves, they have been shown to be a unique type of predicate. Their context-update effect is entirely metalinguistic. They provide interlocutors with a means of speaking directly about their conversation, which allows for the synchronization of the Common Ground in a discourse.

References

Barker, Chris: 2002, 'The Dynamics of Vagueness', *Linguistics and Philosophy*, **25**, 1-36.

Barker, Chris and Gina Taranto: 2002, 'The Paradox of Asserting Clarity', *Proceedings of the Western Conference on Linguistics*, Held Nov. 1-3, University of British Columbia.

Beaver, David: 2002, *Presupposition*, CSLI Publications, Stanford.

Fine, Kit: 1975, 'Vagueness, Truth, and Logic', *Synthese*, **30**, 265-300.

Grice, H.P: 1975, 'Logic and Conversation' in P. Cole and J. Morgan (eds.) *Syntax and Semantics Volume 3: Speech Acts,* Academic Press, New York.

Gunlogson, Christine: 2001, *True to Form: Rising and Falling Declaratives as Questions in English*, PhD Dissertation, University of California, Santa Cruz.

Heim, Irene: 1982, *The Semantics of Definite and Indefinite Noun Phrases.* PhD Dissertation, University of Massachusetts, Amherst.

Heim, Irene: 1983, 'On the Projection Problem for Presuppositions', in D. Flickinger, et al., (eds.), *Proceedings of the Second West Coast Conference on Formal Linguistics*, Stanford University Press, Stanford.

Kiparsky, Paul and Carol Kiparsky: 1970, 'Fact', in M. Bierwisch and K. Heidolph, (ed.s), *Progress in Linguistics*, Mouton, The Hague.

Kyburg, Alice and Michael Moreau: 2000, 'Fitting Words: Vague Language in Context', *Linguistics and Philosophy* **23**, 577-597.

Lewis, David: 1970, 'General Semantics', *Synthese* **22**, 18-67.

van der Sandt, Rob, 1992: Presupposition Projection as Anaphora Resolution', *Journal of Semantics* **9**, 333-377.

Stalnaker, Robert, 1979: 'Assertion', in P. Cole (ed.) *Pragmatics, Syntax and Semantics* 9, Academic Press, New York.

Stalnaker, Robert, 1998: 'On the Representation of Context', *Journal of Logic, Language, and Information,* **7**, 3-19.

Taranto, Gina, 2003: *Discourse Adjectives*, PhD Dissertation, University of California, San Diego.

Walker, Marilyn, 1992: 'Redundancy in Collaborative Dialogue'. *Proceedings of COLING 1992.*

Walker, Marilyn, and Steve Whittaker, 1990: 'Mixed Initiative in Dialogue: An Investigation into Discourse Segmentation'. *The Proceedings of the 28[th] Annual Meeting of the ACL.*

Williamson, Timothy, 1994: *Vagueness*, Routledge, London.

Williamson, Timothy, 1999: 'On the Structure of Higher-Order Vagueness'. *Mind* **108**,127-114.